When Do Programmable Vouchers Improve Public Spending?

An evaluation model for rule enforcement, access, and administrative cost

Jace Shim · November 2026

Abstract

Programmable vouchers can verify conditions such as merchant eligibility and expiry when a beneficiary attempts payment. This paper asks when enforcing those conditions at checkout improves public spending relative to an existing rule-based voucher. It develops an evaluation model that tracks whether eligible recipients complete intended purchases, whether ineligible purchases are prevented, and how the intervention shifts costs across agencies, merchants, and households. A hypothetical example shows why fewer improper payments can coincide with poorer program performance. Korea's Project Hangang and Singapore's Purpose Bound Money work provide implementation context. Three propositions and a pilot design specify how the model could be tested.

1. Research question and contribution

Chapters 6 and 12 of Dollar on Chain propose local digital vouchers with rules that execute during payment (Shim, 2026). A municipal agency might limit a food benefit to designated merchants and a particular period. The same rule can be checked by merchants, audited later, or enforced by a payment system before settlement. These designs distribute error, delay, and administrative work differently.

The contribution is a comparative evaluation framework for deciding when to move enforcement to checkout. It combines enforcement accuracy, benefit delivery, and implementation burden in one decision. The central proposition is that an apparent improvement in compliance can conceal a loss of access when a system blocks eligible purchases or leads households to give up before checkout.

The unit of analysis is the eligible household across a benefit cycle. A secondary unit is the transaction attempt. The policy comparison holds eligibility, benefit value, expiry, and merchant rules fixed wherever possible. The relevant baseline is the program's actual processing arrangement, which may already automate some checks.

2. Evaluation model

Register three outcomes per 1,000 eligible households in each design over the same observation window. T is the number of households that complete a policy-eligible purchase, with a prespecified threshold when partial redemption is common. D is the expected harm of noncompliant spending that succeeds, assessed under a common audit protocol. C is the total implementation and use cost across the program agency, merchants, and households. Cost includes integration and operation, reconciliation, training, appeal handling, and time and travel. Record who incurs each cost and avoid counting pure transfers as resource costs.

The change in assessed program value from introducing checkout enforcement is:

Change in value = w × change in T − change in D − change in C.

Here w is the chosen incremental social value of reaching an additional eligible household, measured in the same monetary units as D and C. Specify benefit valuation and what counts as harm before the evaluation. Use a fixed time horizon and amortize setup costs across scenarios. Report outcomes separately and perform sensitivity analysis for w and for the assumed cost of improperly directed spending.

The model adds a distributional constraint: set a minimum acceptable eligible completion rate and a maximum acceptable decline for groups with access barriers before deployment. Compare these outcomes by merchant density, device access, and need for assisted transactions. A favorable aggregate result can coexist with a subgroup breach. Include unsuccessful attempts and beneficiaries who never reach checkout in the analysis.

3. Mechanism and hypothetical example

Rules based on an authoritative merchant registry and reliable expiry timestamps may be straightforward to apply at checkout. Rules depending on a merchant's changing inventory, ambiguous item categories, or hardship circumstances can be harder to encode. Outdated records and system failures may block authorized purchases. A staffed override or later review can recover such transactions while preserving an auditable decision trail.

Suppose checkout enforcement reduces improper-spending harm and total administrative burden by a combined 9,000 monetary units per 1,000 eligible households. If 60 fewer households complete eligible purchases and each lost completion is valued at 200 units, assessed value declines by 3,000. The example uses invented quantities. It illustrates the break-even value: 150 units per lost completion. A subgroup access breach can independently trigger a redesign under the predeclared minimum standard.

4. Testable propositions

P1. Rule verifiability. The reduction in successful noncompliant purchases should be larger for rules tied to reliable, machine-readable records than for rules requiring context at checkout, conditional on baseline compliance and merchant coverage.

P2. Recoverability and access. A documented assisted path and a time-bounded override should reduce the eligible completion gap associated with checkout enforcement. Test the

effect separately where merchant coverage or device access is limited. Track overrides, subsequent authorization, and time to resolution.

P3. Distribution of burden. Automated checks should reduce manual reconciliation when programs can reuse stable payment integrations. In fragmented settings, merchant onboarding, software maintenance, and beneficiary support may offset those savings. Measure total costs and report them separately by actor.

5. Case context and pilot design

The Bank of Korea describes Project Hangang as an experiment combining wholesale central bank digital currency with bank-issued tokenized deposits. Its phase-one reporting includes a digital voucher program. The Monetary Authority of Singapore's Purpose Bound Money whitepaper describes how conditions can be attached to use while preserving the underlying payment instrument. These examples motivate design comparisons. Evidence of technical operation alone leaves causal effects on access, targeting, and administrative costs open.

For a prospective test, compare three versions of the same benefit: existing checks, checkout enforcement, and checkout enforcement with an assisted exception path. Randomize rollout at a suitable district or merchant-cluster level if operationally feasible, and preserve an alternative means of receiving essential benefits. Pre-register outcomes, subgroup thresholds, and independent adjudication rules. Identify rollout spillovers and estimate effects by assigned treatment. Where randomization is unavailable, use a phased evaluation with credible comparison areas and inspect pre-pilot trends.

Link privacy-minimized administrative records to sampled audits, time-use measures, and surveys of non-redeemers. Attribute failures to the rule, merchant coverage, connectivity, or procedural barriers when the evidence allows. Compare attempted denials with independently adjudicated eligibility and separate household-level nonuse from repeated transaction attempts. Record exceptions and appeals without collecting unnecessary purchase detail.

6. Implications

The framework gives public agencies a decision rule for whether to enforce a condition before payment and when to retain human review. The model also applies to ordinary digital payment systems that can enforce the same rules. Comparing equivalent designs helps identify whether any benefit comes from the payment architecture, better registries, or a redesigned administrative process. This paper specifies hypotheses and evaluation methods. It presents no causal estimates.

References

Bank of Korea. (2026). Monetary Policy Report, September 2026. https://www.bok.or.kr/portal/bbs/B0000156/view.do?menuNo=200067&nttId=11064613

Monetary Authority of Singapore. (2023). Purpose Bound Money Technical Whitepaper. https://www.mas.gov.sg/publications/monographs-or-information-paper/2023/purpose-bound-money-whitepaper

Shim, J. (2026). Dollar on Chain. Raon Book. ISBN 979-11-998906-2-6.